

The Parameterized Complexity of some Permutation Group Problems

V. Arvind
 The Institute of Mathematical Sciences
 C.I.T. Campus
 Chennai 600 113, India
 arvind@imsc.res.in

Abstract

In this paper we study the parameterized complexity of two well-known permutation group problems which are NP-complete.

- Given a permutation group $G = \langle S \rangle \leq S_n$ and a parameter k , find a permutation $\pi \in G$ such that $|\{i \in [n] \mid \pi(i) \neq i\}| \geq k$. This generalizes the NP-complete problem of finding a fixed-point free permutation in G [CW10, Lub81] (this is the case when $k = n$). We show that this problem with parameter k is fixed parameter tractable. In the process, we give a simple deterministic polynomial-time algorithm for finding a fixed point free element in a transitive permutation group, answering an open question of Cameron [C11, CW10].
- Next we consider the problem of computing a base for a permutation group $G = \langle S \rangle \leq S_n$. A *base* for G is a subset $B \subseteq [n]$ such that the subgroup of G that fixes B pointwise is trivial. This problem is known to be NP-complete [Bl92]. We show that it is fixed parameter tractable for the case of cyclic permutation groups and for permutation groups of constant orbit size. For more general classes of permutation groups we do not know whether the problem is in FPT or is W[1]-hard.

1 Introduction

Let S_n denote the group of all permutations on a set of size n . The group S_n is also called the symmetric group of degree n . We refer to a subgroup G of S_n , denoted by $G \leq S_n$, as a permutation group (of degree n). Let $S \subseteq S_n$ be a subset of permutations. The permutation group *generated* by S , denoted by $\langle S \rangle$, is the smallest subgroup of S_n containing S . A subset $S \subseteq G$ of a permutation group G is a *generating set* for G if $G = \langle S \rangle$. It is easy to see that every finite group G has a generating set of size $\log_2 |G|$.

Let $G = \langle S \rangle \leq S_n$ be a subgroup of the symmetric group S_n , where G is given as input by a generating set S of permutations. There are many algorithmic problems on permutation groups that are given as input by their generating

sets (e.g. see [Sim70, FHL80, Luk93, Ser03]). Some of them have efficient algorithms, some others are NP-complete, and yet others have a status similar to Graph Isomorphism: they are neither known to be in polynomial time and unlikely to be NP-complete (unless the Polynomial-Time Hierarchy collapses). Efficient permutation group algorithms have played an important role in the design of algorithms for the Graph Isomorphism problem [Bab79, BKL83]. In fact the algorithm with the best running time bound for general Graph Isomorphism is group-theoretic.

We recall some definitions and notions from permutation group theory. Let $\pi \in S_n$ be a permutation. A *fixed point* of π is a point $i \in [n]$ such that $\pi(i) = i$ and π is *fixed point free* if $\pi(i) \neq i$ for all $i \in [n]$.

Let $G \leq S_n$ and $\Delta \subseteq [n]$ be a subset of the domain. The *pointwise stabilizer subgroup* of G , denoted G_Δ , is $\{g \in G \mid g(i) = i \text{ for all } i \in \Delta\}$.

A subset $B \subseteq [n]$ is called a *base* for G if the pointwise stabilizer subgroup G_B is trivial. Thus, if B is a base for G then each element of G is uniquely determined by its action on B . The problem of computing a base of minimum cardinality is known to be computationally very useful. Important algorithmic problems on permutation groups, like membership testing, have nearly linear time algorithms in the case of small-base groups (e.g. see [Ser03]). We will discuss the parameterized complexity of the minimum base problem in Section 3.

An excellent modern reference on permutation groups is Cameron's book [C99]. Algorithmic permutation group problems are very well treated in [Luk93, Ser03]. Basic definitions and results on parameterized complexity can be found in Downey and Fellows' classic text on the subject [DF99]. Another, more recent, reference is [FG06].

2 Fixed point free elements

The starting point is the Orbit-Counting lemma. Our discussion will follow Cameron's book [C99]. For each permutation $g \in S_n$ let $\text{fix}(g)$ denote the number of points fixed by g . More precisely,

$$\text{fix}(g) = |\{i \in [n] \mid g(i) = i\}|.$$

A permutation group $G \leq S_n$ induces, by its action an equivalence relation on the domain $[n]$: i and j are in the same equivalence class if $g(i) = j$ for some $g \in G$. Each equivalence class is an *orbit* of G . G is said to be *transitive* if there is exactly one G -orbit. Let $\text{orb}(G)$ denote the number of G -orbits in the domain $[n]$. We recall the statement.

Lemma 2.1 (Orbit Counting Lemma). [C92] *Let $G \leq S_n$ be a permutation group. Then*

$$\text{orb}(G) = \frac{1}{|G|} \sum_{g \in G} \text{fix}(g). \quad (1)$$

I.e. the number of G orbits is the average number of fixed points over all elements of G .

Proof. It is useful to recall a proof sketch. Define a $|G| \times n$ matrix with rows indexed by elements of G and columns by points in $[n]$. The $(g, i)^{th}$ entry is defined to be 1 if $g(i) = i$ and 0 otherwise. Clearly, the g^{th} row has $\text{fix}(g)$ many 1's in it. Let G_i denote the subgroup of G that fixes i . The i^{th} column clearly has $|G_i|$ many 1's. Counting the number of 1's in the rows and columns and equating them, keeping in mind that $|G|/|G_i|$ is the size of the orbit containing i yields the lemma. \square

We now recall a theorem of Jordan on permutation groups [J72]. See [Se03, C11] for very interesting accounts of it. A permutation group $G \leq S_n$ is *transitive* if it has exactly one orbit.

Theorem 2.2 (Jordan's theorem). *If $G \leq S_n$ is transitive then G has a fixed point free element.*

It follows directly from the Orbit counting lemma. Notice that the left side of Equation 1 equals 1. The right side of the equation is the average over all $\text{fix}(g)$. Now, the identity element 1 fixes all n elements. Thus there is at least one element $g \in G$ such that $\text{fix}(g) = 0$. Cameron and Cohen [C92] do a more careful counting and show the following strengthening.

Theorem 2.3. [C92] *If $G \leq S_n$ is transitive then there are at least $|G|/n$ elements of G that are fixed point free.*

We discuss their proof, because we will build on it to obtain our results. If G is transitive, the orbit counting lemma implies

$$|G| = \sum_{g \in G} \text{fix}(g).$$

Take any point $\alpha \in [n]$. We can write the above equation as

$$|G| = \sum_{g \in G_\alpha} \text{fix}(g) + \sum_{g \in G \setminus G_\alpha} \text{fix}(g).$$

By the orbit counting lemma applied to the group G_α we have

$$\sum_{g \in G_\alpha} \text{fix}(g) = \text{orb}(G_\alpha) \cdot |G_\alpha|.$$

Let $F \subset G$ be the set of all fixed point free elements of G . Clearly, $\sum_{g \in G \setminus G_\alpha} \text{fix}(g) \geq |G \setminus A|$ as $A \subseteq G \setminus G_\alpha$ and each element of $G \setminus A$ fixes at least one element. Combining with the previous equation we get

$$|A| \geq \text{orb}(G_\alpha) \cdot |G_\alpha| = \text{orb}(G_\alpha) \cdot \frac{|G|}{n} \geq \frac{|G|}{n}.$$

2.1 The Algorithmic Problem

We now turn to the problem of computing a fixed point free element in a permutation group $G \leq S_n$ and a natural parameterized version.

As observed by Cameron and Wu in [CW10], the result of [C92] gives a simple randomized algorithm to find a fixed point free element in a transitive permutation group $G \leq S_n$, where G is given by a generating set S : Using Schreier-Sims polynomial-time algorithm [Sim70] we can compute a *strong generating set* S' for G in polynomial time. And using S' we can sample uniformly at random from G . Clearly, in $O(n)$ sampling trials we will succeed in finding a fixed point free element with constant probability. We will show in the next section that this algorithm can be *derandomized* to obtain a deterministic polynomial time algorithm (without using CFSG). This answers an open problem of Cameron discussed in [CW10, C11].

This result is to be contrasted with the fact that computing fixed point free elements in nontransitive groups $G \leq S_n$ is NP-hard. The decision problem is shown NP-complete in [CW10]. This is quite similar to Lubiw's result [Lub81] that checking if a graph X has a fixed point free automorphism is NP-complete.

We will now introduce the parameterized version of the problem of computing fixed point free elements in permutation groups. First we introduce some terminology. We say that a permutation π *moves* a point $i \in [n]$ if $\pi(i) \neq i$.

k -MOVE Problem

INPUT: A permutation group $G = \langle S \rangle \leq S_n$ given by generators and a number k .

PROBLEM: Is there an element $g \in G$ that moves at least k points.

For $k = n$ notice that k -MOVE is precisely the problem of checking if there is a fixed point free element in G . The parameterized version of the problem is to treat k as parameter. We will show that this problem is fixed parameter tractable.

Let $\text{move}(g)$ denote the number of points moved by g . We define two numbers $\text{fix}(G)$ and $\text{move}(G)$:

$$\begin{aligned} \text{fix}(G) &= |\{i \in [n] \mid g(i) = i \text{ for all } g \in G\}| \\ \text{move}(G) &= |\{i \in [n] \mid g(i) \neq i \text{ for some } g \in G\}| \end{aligned}$$

I.e. $\text{fix}(G)$ is the number of points fixed by all of G and $\text{move}(G)$ is the number of points moved by some element of G . Clearly, for all $g \in G$, $\text{move}(g) = n - \text{fix}(g)$ and $\text{move}(G) = n - \text{fix}(G)$. Furthermore, notice that $\text{orb}(G) \leq \text{fix}(G) + \text{move}(G)/2$, and we have $n - \text{orb}(G) \geq \text{move}(G)/2$. Let $G = \langle S \rangle \leq S_n$ be an input instance for the k -MOVE problem. Substituting $n - \text{move}(g)$ for

$\text{fix}(g)$ in Equation 1 and rearranging terms we obtain

$$\text{move}(G)/2 \leq n - \text{orb}(G) = \frac{1}{|G|} \sum_{g \in G} \text{move}(g) = \mathbb{E}_{g \in G}[\text{move}(g)], \quad (2)$$

where the expectation is computed for g picked uniformly at random from G .

We will show there is a deterministic polynomial time algorithm that on input $G = \langle S \rangle \leq S_n$ outputs a permutation $g \in G$ such that $\text{move}(g) \geq n - \text{orb}(G) \geq \text{move}(G)/2$. Using this algorithm we will obtain an FPT algorithm for the k -MOVE problem. We require the following useful lemma about computing the average number of points moved by uniformly distributed elements from a coset contained in S_n .

Lemma 2.4. *Let $G\pi \subseteq S_n$ be a coset of a permutation group $G = \langle S \rangle \leq S_n$, where $\pi \in S_n$. There is a deterministic algorithm that computes $\mathbb{E}_{g \in G}[\text{move}(g\pi)]$ in time polynomial in $|S|$ and n .*

Proof. We again use a double counting argument. Define a 0-1 matrix with rows indexed by $g\pi, g \in G$ and columns by $i \in [n]$, whose $(g\pi, i)^{\text{th}}$ entry is 1 if and only if $g(\pi(i)) \neq i$. Thus, the number of 1's in the i^{th} column of the matrix is $|G| - |\{g \in G \mid g(\pi(i)) = i\}|$. Now, $|\{g \in G \mid g(\pi(i)) = i\}|$ is zero if $\pi(i)$ and i are in different G -orbits and is $|G_i|$ if they are in the same orbit. In polynomial time we can compute the orbits of G and check this condition. Also, the number $|G| - |\{g \in G \mid g(\pi(i)) = i\}| = |G| - |G_i|$ is computable in polynomial time. Call this number N_i . It follows that the total number of 1's in the matrix is $\sum_{i=1}^n N_i$, which is computable in polynomial time. Since $\sum_{i=1}^n N_i = \sum_{g \in G} \text{move}(g\pi)$, it follows that $\frac{1}{|G|} \sum_{g \in G} \text{move}(g\pi) = \mathbb{E}_{g \in G}[\text{move}(g\pi)]$ can be computed exactly in polynomial time. \square

Theorem 2.5. *There is a deterministic polynomial-time algorithm that takes as input a permutation group $G = \langle S \rangle \leq S_n$ given by generators and a permutation $\pi \in S_n$ and computes an element $g \in G$ such that $\text{move}(g\pi) \geq \mathbb{E}_{g \in G}[\text{move}(g\pi)]$.*

Proof. We have

$$\frac{1}{|G|} \sum_{g \in G} \text{move}(g\pi) = \mathbb{E}_{g \in G}[\text{move}(g\pi)] = \mu,$$

and by Lemma 2.4 we can compute μ in polynomial time. We can write G as a disjoint union of cosets $G = \bigcup_{i=1}^r G_1 g_i$, where G_1 is the subgroup of G that fixes 1 and g_i are the coset representatives, where the number of cosets $r \leq n$. Using Schreier-Sims algorithm [Sim70] we can compute all coset representatives g_i and a generating set for G_1 from the input in polynomial time.

Now, we can write the summation $\frac{1}{|G|} \sum_{g \in G} \text{move}(g\pi)$ as a sum over the cosets $G_1 g_i \pi$ of G_1 :

$$\frac{1}{|G|} \sum_{g \in G} \text{move}(g\pi) = \frac{1}{|G|} \sum_{i=1}^r \sum_{g \in G_1} \text{move}(gg_i\pi).$$

For $1 \leq i \leq r$ let

$$\mu_i = \frac{1}{|G_1|} \sum_{g \in G_1} \text{move}(gg_i\pi).$$

Since $|G|/|G_1| = r$, it follows that $\mu = \frac{1}{r} \sum_{i=1}^r \mu_i$ is an average of the μ_i . Let μ_t denote $\max_{1 \leq i \leq r} \mu_i$. Clearly, $\mu \leq \mu_t$ and therefore there is some $g \in G_1 g_t \pi$ such that $\text{move}(g) \geq \mu_t \geq \mu$ and we can continue the search in the coset $G_1 g_t$ since we can compute all the μ_i in polynomial time by Lemma 2.4. Continuing thus for $n-1$ steps, in polynomial time we will obtain a coset $G_{n-1} \tau$ containing the unique element τ such that $\text{move}(\tau) \geq \mu$. This completes the proof. \square

Cameron, in [CW10] and in the lecture notes [C11], raises the question whether the randomized algorithm, based on uniform sampling, for finding a fixed point free element in a transitive permutation group (given by generators) can be derandomized. In [CW10] a deterministic algorithm (based on the classification of finite simple groups) is outlined. The algorithm does a detailed case analysis based on the CFSG and is not easy to verify. Here we show that the randomized algorithm can be easily derandomized yielding a simple polynomial-time algorithm. The derandomization is essentially a simple application of the “method of conditional probabilities” [ES73, Ra88].

Corollary 2.6. *Given a transitive permutation group $G = \langle S \rangle \leq S_n$ by a generating set S , we can compute a fixed point free element of G in deterministic polynomial time.*

Proof. Notice that $\mathbb{E}_{g \in G}[\text{move}(g)] = n-1$ to begin with. However, since G_1 has at least two orbits, we have by orbit counting lemma that $\mathbb{E}_{g \in G_1}[\text{move}(g)] \leq n-2$. Hence, for some coset $G_1 g_i$ of G_1 in G we must have $\mathbb{E}_{g \in G_1 g_i}[\text{move}(gg_i)] > n-1$. The polynomial-time algorithm of Theorem 2.5 applied to G will therefore continue the search in cosets where the expected value is strictly more than $n-1$ which means that it will finally compute a fixed point free element of G . \square

Given $G = \langle S \rangle \leq S_n$ there is a trivial exponential time algorithm for finding a fixed point free element in G : compute a strong generating set for G in polynomial time [Sim70]. Then enumerate G in time $|G| \cdot n^{O(1)}$ using the strong generating set, checking for a fixed point free element. This algorithm could have running time $n!$ for large G . We next describe a $2^n n^{O(1)}$ time algorithm for finding a fixed point free element based on inclusion-exclusion and coset intersection.

Theorem 2.7. *Given a permutation group $G = \langle S \rangle \leq S_n$ and $\pi \in S_n$ there is a $2^{n+O(\sqrt{n} \lg n)} n^{O(1)}$ time algorithm to test if the coset $G\pi$ has a fixed point free element and if so compute it.*

Proof. For each subset $\Delta \subseteq [n]$ we can compute the pointwise stabilizer subgroup G_Δ . This will take time $2^n n^{O(1)}$ overall. For each $i \in [n]$, let $(G\pi)_i$ denote the subcoset of $G\pi$ that fixes i . Indeed,

$$(G\pi)_i = \{g\pi \mid g \in G, g\pi(i) = i\} = G_{\pi(i)} \tau_i \pi,$$

if there is a $\tau_i \in G$ such that $\tau_i(\pi(i)) = i$ and $(G\pi)_i = \emptyset$ otherwise.

Clearly, $G\pi$ has a fixed point free element if and only if the union $\bigcup_{i=1}^n (G\pi)_i$ is a *proper* subset of $G\pi$. I.e. we need to check if $|\bigcup_{i=1}^n (G\pi)_i| < |G\pi| = |G|$. Now, $|\bigcup_{i=1}^n (G\pi)_i|$ can be computed in $2^{n+O(\sqrt{n} \lg n)} n^{O(1)}$ time using the inclusion exclusion principle: there are 2^n terms in the inclusion-exclusion formula. Each term is the cardinality of a coset intersection of the form $\bigcap_{i \in I} (G\pi)_i$, for some subset of indices $I \subseteq [n]$, which can be computed in time $n^{O(\sqrt{n})}$ time [BKL83]. Hence, we can decide in $2^{n+O(\sqrt{n} \lg n)} n^{O(1)}$ time whether or not $G\pi$ has a fixed point free element. Notice that this fixed point free element must be in one of the $n - 1$ subcosets of $G\pi$ that maps 1 to j for $j \in \{2, 3, \dots, n\}$. The subcoset of $G\pi$ mapping 1 to j can be computed in polynomial time [Sim70]. Then we can apply the inclusion exclusion principle to each of these subcosets, as explained above, to check if it contains a fixed point free element and continue the search in such a subcoset. Proceeding thus for $n - 1$ steps we will obtain a fixed point free element in $G\pi$, if it exists, in $2^{n+O(\sqrt{n} \lg n)} n^{O(1)}$ time. \square

We now prove the main result of this section.

Theorem 2.8. *There is a deterministic $2^{2k+O(\sqrt{k} \lg k)} k^{O(1)} + n^{O(1)}$ time algorithm for the k -MOVE problem and hence the problem is fixed parameter tractable. Furthermore, if $G = \langle S \rangle \leq S_n$ is a “yes” instance the algorithm computes a $g \in G$ such that $\text{move}(g) \geq k$.*

Proof. Let $G = \langle S \rangle \leq S_n$ be an input instance of k -MOVE with parameter k . By Equation 2 we know that $\mathbb{E}_{g \in G} [\text{move}(g)] \geq \text{move}(G)/2$. We first compute $\text{move}(G)$ in polynomial time by computing the orbits of G . If $\text{move}(G) \geq 2k$ then the input is a “yes” instance to the problem and we can apply Theorem 2.5 to compute a $g \in G$ such that $\text{move}(g) \geq k$ in polynomial time. Otherwise, $\text{move}(G) \leq 2k$. In that case, the group G is effectively a permutation group on a set $\Omega \subseteq [n]$ of size at most $2k$. For each subset $\Delta \subseteq \Omega$ of size at most k , we compute the pointwise stabilizer subgroup G_Δ of G in polynomial time [Sim70]. This will take overall $2^{2k} n^{O(1)}$ time. Now, if the input is a “yes” instance to k -MOVE, some subgroup G_Δ must contain a fixed point free element (i.e. fixed point free in $\Omega \setminus \Delta$). We can apply the algorithm of Theorem 2.7 to compute this element in time $2^{2k+O(\sqrt{k} \lg k)} k^{O(1)}$. \square

Remark. We note from the first few lines in the proof of Theorem 2.8 that the application of Theorem 2.5 is actually a polynomial time reduction from the given k -MOVE instance to an instance for which $\text{move}(G) \leq 2k$. Given $G = \langle S \rangle \leq S_n$ such that $\text{move}(G) \leq 2k$, note that G is effectively a subgroup of S_{2k} . We can apply the Schreier-Sims algorithm to compute from S a generating set of size $O(k^2)$ for G , therefore yielding a polynomial time computable, $k^{O(1)}$ size kernel (see [FG06] for definition) for the k -MOVE problem.

3 The parameterized minimum base problem

In this section we turn to another basic algorithmic problem on permutation groups.

Definition 3.1. Let $G \leq S_n$ be a permutation group. A subset of points $B \subseteq [n]$ is called a *base* if the pointwise stabilizer subgroup G_B of G (subgroup of G that fixes B pointwise) is the identity.

Since permutation groups with a small base have fast algorithms for various problems [Ser03], computing a minimum cardinality base for G is very useful. The decision problem is NP-complete. On the other hand, it has a $\lg \lg n$ factor approximation algorithm [Bl92].

In this section we study the parameterized version of the problem with base size as parameter. We are unable to resolve if the general case is FPT or not, we give FPT algorithms in the case of cyclic permutation groups and for permutation groups with orbits of size bounded by a constant.

k -BASE Problem

INPUT: A permutation group $G = \langle S \rangle \leq S_n$ given by generators and a number k .

PROBLEM: Is there a base of size at most k for G . The search version is the find such a base.

A trivial $n^{k+O(1)}$ algorithm would cycle through all candidate subsets B of size at most k checking if G_B is the identity.

Remark. If the elements of the group $G \leq S_n$ are explicitly listed, then the k -BASE problem is essentially a hitting set problem, where the hitting set B has to intersect, for each $g \in G$, the subset of points moved by g . However, the group structure makes it different from the general hitting set problem and we do not know how to exploit it algorithmically in the general case.

3.1 Cyclic Permutation Groups

We give an FPT algorithm for the special case when the input permutation group $G = \langle S \rangle$ is cyclic. While this is only a special case, we note that the minimum base problem is NP-hard even for cyclic permutation groups [Bl92, Theorem 3.1].

Theorem 3.2. *The k -BASE problem for cyclic permutation groups is fixed parameter tractable.*

Proof. Let $G = \langle S \rangle \leq S_n$ be a cyclic permutation group as instance for k -BASE. Using known polynomial-time algorithms [Sim70, Luk93] we can compute a decomposition of G into a direct product of cyclic groups of prime power order.

$$G = H_1 \times H_2 \times \dots \times H_\ell$$

where each H_i is cyclic of prime power order. Let $H_i = \langle g_i \rangle$, where the order of g_i , $o(g_i) = p_i^{e_i}$, $1 \leq i \leq \ell$, where the p_i 's are all distinct. Notice that

$|G| = p_1^{e_1} p_2^{e_2} \dots p_\ell^{e_\ell}$. We can assume $|G| \leq n^k$. Otherwise, G does not have a size k base and the algorithm can reject the instance. Since

$$(\ell/e)^\ell \leq \ell! \leq p_1 p_2 \dots p_\ell \leq n^k,$$

it follows that $\ell = O(\frac{k \lg n}{\lg \lg n})$.

For each g_i , when we express it as a product of disjoint cycles then the length of each such cycle is a power of p_i that divides $p_i^{e_i}$, and there is at least one cycle of length $p_i^{e_i}$. Clearly, any base for G must include at least one point of some $p_i^{e_i}$ -cycle (i.e. cycle of length $p_i^{e_i}$) of g_i , for each i . Otherwise, the cyclic subgroup H_i of G will not become identity when the points in the base are fixed. For each index $i : 1 \leq i \leq \ell$, define the set of points

$$S_i = \{\alpha \in [n] \mid \alpha \text{ is in some } p_i^{e_i} \text{ cycle of } g_i\}.$$

Claim. Let $B \subseteq [n]$ be a subset of size k . Then B is a base for G if and only if B is a hitting set for the collection of sets $\{S_1, S_2, \dots, S_\ell\}$.

Proof of Claim. Clearly, it is a necessary condition. Conversely, suppose $|B| = k$ and $B \cap S_i \neq \emptyset$ for each i . Consider the partition of $[n]$ into the orbits of G :

$$[n] = \Omega_1 \cup \Omega_2 \cup \dots \cup \Omega_r.$$

For each g_i , a cycle of length $p_i^{e_i}$ in g_i is wholly contained in some orbit of G . Indeed, each orbit of G must be a union of a subset of cycles of g_i . Since $B \cap S_i \neq \emptyset$, some $p_i^{e_i}$ -cycle C_i of g_i will intersect B .

Assume, contrary to the claim, that there is a $g \in G_B$ such that $g \neq 1$. We can write $g = g_1^{a_1} g_2^{a_2} \dots g_\ell^{a_\ell}$ for nonnegative integers $a_i < p_i^{e_i}$. Suppose $g_j^{a_j} \neq 1$. Then raising both sides of the equation $g = g_1^{a_1} g_2^{a_2} \dots g_\ell^{a_\ell}$ to the power $\frac{|G|}{p_j^{e_j}}$, we have

$$g' = g^{\frac{|G|}{p_j^{e_j}}} = g_j^{\beta_j},$$

where $\beta_j < p_j^{e_j}$. Moreover, $\beta_j = \frac{|G| a_j}{p_j^{e_j}} \pmod{p_j^{e_j}}$ is nonzero because $a_j \neq 0 \pmod{p_j^{e_j}}$ and $|G|/p_j^{e_j}$ does not have p_j as factor.

By assumption, some $p_j^{e_j}$ -cycle C_j of g_j intersects B . Since β_j is nonzero and strictly smaller than $p_j^{e_j}$, none of the points of C_j are fixed by $g_j^{\beta_j}$ which contradicts the assumption that g and hence g' is in G_B . This proves the claim.

We now explain the FPT algorithm. If $|G| > n^k$ then there is no base of size k . Hence we can assume $|G| \leq n^k$. As already observed, $\ell = O(\frac{k \lg n}{\lg \lg n})$. Thus, we need to solve the k -hitting set problem for a collection of at most $O(\frac{k \lg n}{\lg \lg n})$ many sets $\{S_1, S_2, \dots, S_\ell\}$. We can think of it as a problem of k -coloring the indices $\{1, 2, \dots, \ell\}$ such that for each color class I we have $\cap_{i \in I} S_i \neq \emptyset$ and we can pick any one point for each such intersection. Notice that there are at most $k^\ell = n^{\frac{k \lg k}{\lg \lg n}}$ many such colorings. Now, if $k \lg k \leq \lg \lg n$ this number is bounded by $n^{O(1)}$ and we can cycle through all these k -colorings in polynomial time and

find a good k -coloring if it exists. On the other hand, if $k \lg k > \lg \lg n$ then $n^k \leq 2^{k^{k+1}}$ which means the brute force search gives an FPT time bound. \square

3.2 Bounded Orbit Permutation Groups

We give an FPT algorithm for another special case of the k -BASE problem: Let $G = \langle S \rangle \leq S_n$ such that G has orbits of size bounded by a fixed constant b . I.e. $[n] = \bigsqcup_{i=1}^m \Omega_i$, where $|\Omega_i| \leq b$ for each i . This is again an interesting special case as the minimum base problem is NP-hard even for orbits of size bounded by 8 [Bl92, Theorem 3.2].

Suppose G has a base $B = \{i_1, i_2, \dots, i_k\}$ of size k . Then G has a pointwise stabilizer tower $G = G_0 \geq G_1 \geq \dots \geq G_k = \{1\}$ obtained by successively fixing the points of B . More precisely, G_j is the subgroup of G that pointwise fixes $\{i_1, i_2, \dots, i_j\}$. Now, $\frac{|G_{j-1}|}{|G_j|}$ is the orbit size of the point i_j in the group G_{j-1} . Furthermore, b is also a bound on this orbit size. Therefore, $|G| \leq b^k$. Hence in $b^k n^{O(1)}$ time we can list all elements of G . Let $G = \{g_1, g_2, \dots, g_N\}$, where $N \leq b^k$, where g_1 is the identity element.

For each $g_i \in G, i \geq 2$, let $S_i = \{j \in [n] \mid g_i(j) \neq j\}$ denote the nonempty subset of points not fixed by g_i . Then a subset $B \subset [n]$ of size k is a base for G if and only if B is a hitting set for the collection S_2, S_3, \dots, S_N . The next claim is straightforward.

Claim. There is a size k hitting set contained in $[n]$ for the sets $\{S_2, S_3, \dots, S_N\}$ if and only if there is a partition of $\{2, 3, \dots, N\}$ into k parts I_1, I_2, \dots, I_k such that $\cap_{j \in I_r} S_j \neq \emptyset$ for each $r = 1, 2, \dots, k$.

As $N \leq b^k$, the total number of k -partitions of $\{2, 3, \dots, N\}$ is bounded by $k^N \leq k^{b^k}$. We can generate them and check if any one of them yields a hitting set of size k by checking the condition in the above claim. The overall time taken by the algorithm is given by the FPT time bound $k^{b^k} n^{O(1)}$. We have shown the following result.

Theorem 3.3. *Let $G = \langle S \rangle \leq S_n$ such that G has orbits of size bounded by b , be an instance for the k -BASE problem with k as parameter. Then the problem has an FPT algorithm of running time $k^{b^k} n^{O(1)}$.*

4 Concluding Remarks

The impact of parameterized complexity on algorithmic graph theory research, especially its interplay with graph minor theory, has been very fruitful in the last two decades. This motivates the study of parameterized complexity questions in other algorithmic problem domains like, for example, group-theoretic computation. To this end, we considered parameterized versions of two well-known classical problems on permutation groups. We believe that a similar study of other permutation group problems can be a worthwhile direction.

References

- [Bab79] L. Babai. Monte-Carlo algorithms in graph isomorphism testing. Technical Report 79–10, Universit de Montral, 1979.
- [BKL83] L. Babai, W.M. Kantor, and E.M. Luks. Computational complexity and the classification of finite simple groups. In *Proceedings of 24th IEEE Symposium on Foundations of Computer Science (SFCS)*, pages 162–171, 1983.
- [BL83] L. Babai and E.M. Luks. Canonical labeling of graphs. In *Proceedings of 15th Annual ACM Symposium on Theory of Computing (STOC)*, pages 171–183, 1983.
- [Bl92] K.D. Blaha. Minimal bases for permutation groups: the greedy approximation algorithm. *Journal of Algorithms*, 13: 297–306, 1992.
- [C99] P.J. Cameron. *Permutation Groups*. London Mathematical Society, Student Texts 45, Cambridge Univ Press, Cambridge, 1999.
- [C92] P.J. Cameron and A.M. Cohen. On the number of fixed point free elements of a permutation group. *Discrete Mathematics*, 106/107: 135–138, 1992.
- [CW10] P.J. Cameron and T. Wu. The complexity of the weight problem for permutation and matrix groups. *Discrete Mathematics* 310: 408–416, 2010.
- [C11] P.J. Cameron. Lectures on derangements. *Pretty Structures Conference in Paris*, <http://www.lix.polytechnique.fr/Labo/Leo.Liberti/pretty-structures/pdf/pcameron-ps11.pdf>, 2011.
- [DF99] R. G. Downey and M. R. Fellows. *Parameterized Complexity*. Springer, 1999.
- [ES73] P. Erdős and J.L. Selfridge. On a combinatorial game. *Journal of Combinatorial Theory, Series A*, 14(3): 298–301, 1973.
- [FG06] J. Flum and M. Grohe. *Parameterized Complexity Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer, Berlin, 2006.
- [FHL80] M. Furst, J.E. Hopcroft, and E.M. Luks. Polynomial-time algorithms for permutation groups. Technical report, Cornell University, 10 1980.
- [J72] C. Jordan. Recherches sur les substitutions. *J. Math. Pures Appl. (Liouville)*, 17:351–387, 1872.
- [Lub81] A. Lubiw. Some NP-Complete Problems Similar to Graph Isomorphism. *SIAM Journal of Computing* 10(1):11–21, 1981.

- [Luk93] E.M. Luks. Permutation groups and polynomial-time computation. In Larry Finkelstein and William M. Kantor, editors, *Groups and Computation*, number 11 in Discrete Mathematics and Theoretical Computer Science, pages 139–175. American Mathematical Society, 1993.
- [Ra88] P. Raghavan. Probabilistic construction of deterministic algorithms: approximating packing integer programs. *Journal of Computer and System Sciences* 37(2):130143, 1988.
- [Ser03] A. Seress. *Permutation Group Algorithms*. Cambridge University Press, 2003.
- [Se03] J.P. Serre. On a theorem of Jordan. *Bull. Amer. Math. Soc.* 40:429-440, 2003.
- [Sim70] C. C. Sims. Computational methods in the study of permutation groups. In J. Leech, editor, *Computational problems in abstract algebra, Proc. Conf. Oxford, 1967*, pages 169–183. Pergamon Press, 1970.
- [Sim78] C. C. Sims. Some group theoretic algorithms. In A. Dold and B. Eckmann, editors, *Topics in Algebra*, volume 697 of *Lecture Notes in Mathematics*, pages 108–124. Springer, 1978.